

---

## VERITAS Data Acquisition and Analysis Systems

---

Wei Cui,<sup>1</sup> Filip Pizlo,<sup>1</sup> Sam Nolan,<sup>1</sup> John Quinn,<sup>2</sup> Scott Wakely,<sup>3</sup> Ozlem Celik,<sup>4</sup> and Rene Ong<sup>4</sup>

(1) *Department of Physics, Purdue University, West Lafayette, IN 47907, U.S.A.*

(2) *Department of Experimental Physics, University College Dublin, Belfield, Dublin 4, Ireland*

(3) *University of Chicago, Enrico Fermi Institute, Chicago, IL 60637 U.S.A.*

(4) *Department of Physics and Astronomy, University of California, Los Angeles, Los Angeles, CA 90095, U.S.A.*

---

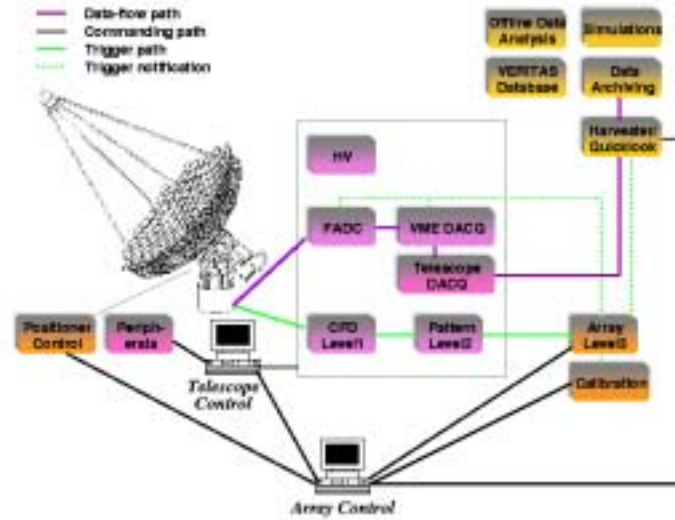
### Abstract

We describe the design of data acquisition and analysis systems for the VERITAS TeV observatory. With an expected average trigger rate as high as 1 kHz, the data acquisition system must be capable of handling data rates up to about 15 MBytes/s. We describe test results on the performance of a prototype system and show that the system has already exceeded the specifications. Data stream from the data acquisition system is fed, in real time, to a data analysis system that provides critical diagnostic information, as well as alerts of bright transient phenomena of scientific interest. Final data products are stored as HDF5 files to facilitate offline analysis. We describe a general framework within which offline analysis software may be developed.

### 1. Introduction

Over the past decade, one of the most exciting advances in high energy astrophysics has been the detection of astronomical objects at TeV energies with ground-based gamma-ray facilities (e.g., review by Catanese & Weekes [1]). Efforts are now being focused on building next-generation TeV observatories that promise to deliver significantly higher sensitivity and lower energy threshold (see [2] for a recent review). The Very Energetic Radiation Imaging Telescope Array System (VERITAS) will be one of such observatories [3]. VERITAS will initially consist of an array of four 12-m telescopes and may, subsequently, be expanded to seven telescopes. A prototype telescope system is near completion (Wakely et al., these proceedings). VERITAS-4 is expected to come online in 2005.

Software development for VERITAS spans over many tasks, from data acquisition, data analysis, and data archiving, to hardware control, calibration, and inter-process communication. Fig. 1 shows a schematic layout of the individual



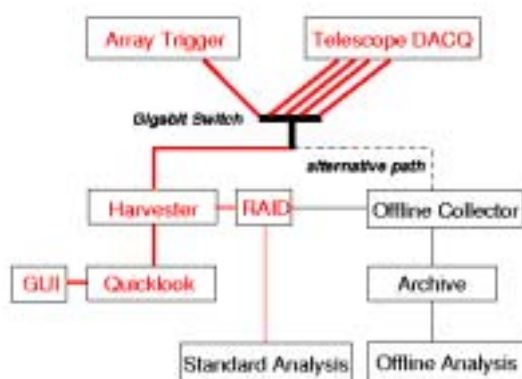
**Fig. 1.** VERITAS software systems. Note that only one telescope is shown.

components. Data acquisition process is the subject of this paper. Some of the other software components are described elsewhere (Kosack et al., Kieda et al., Krawczynski et al., Fegan et al., these proceedings).

## 2. Telescope Data Acquisition

The front-end electronics consist of 5 VME crates, with each being controlled by a Pentium crate computer running the Linux operating system. Each VME computer in a telescope is connected to each other and to the data-acquisition (DACQ) computer via the Scalable Coherent Interface (SCI). The SCI standard describes a packet-based protocol using unidirectional links that provides participating nodes with a shared-memory view of the system. The advantages of using SCI include low latency, deterministic behavior, guaranteed transfer rates, and arbitration handled in hardware.

Upon triggering, each VME computer reads out data from FADC and auxiliary boards in the crate and buffers it in local memory before transferring them in blocks to the DACQ computer, after an optimum number have been accumulated. The advantage of this buffering is that the DACQ computer is largely insulated from fluctuations in the trigger rate and real-time response is not required — it must only be able to keep up with the average event rate. The DACQ system then checks the integrity of the data and puts together all the bits and pieces to form telescope events. The telescope events are buffered and transferred to the array data acquisition system (Harvester). In parallel, they are written locally to the disk and saved for one night. The DACQ system runs on a dual-process Xeon (2.4 GHz) server.



**Fig. 2.** Data acquisition processes. Runtime processes are indicated in red, while offline processes are in black.

With an expected average trigger rate as high as 1 kHz, the DACQ chain must be capable of handling a data rate of about 2.2 MBytes/s. This specification has easily been met. For instance, in a recent data-flow test, the prototype system routinely ran at a rate of nearly 4 MBytes/s. We expect that the limiting factor lies in the VME backplane hardware read-out.

### 3. Array Data Acquisition

The Harvester collects events from individual telescopes that are grouped in a sub-array configuration, checks their integrity, matches them up, and assembles them into array events. The array events are fed to a real-time analysis system (Quicklook) and are also saved locally to the disk. During daytime, the data archiving engine (Offline Collector) writes the array events to HDF5 files and saves the files in the archiving RAID system for offline analyses. Fig. 2 shows a flow chart of the processes.

The Harvester is nearly complete. It was designed for VERITAS-7, which is expected to produce data at rates up to 15 Mbytes/s. Throughput tests have shown that the system can already handle data rates more than twice as high, although the tests was conducted with telescope event simulators and could not have taken into account such factors as latency in a real system. Nevertheless, the results are encouraging. The Offline Collector is being implemented. The purpose of having such a system is to ensure the quality of final data products. The Offline Collector reads in raw array events from the Harvester RAID, matches up “orphan events” (which are incomplete array events caused by, e.g., a significant delay in the arrival of data from one of the telescopes in the sub-array), sort all the events, and save them in HDF5 files. When it is determined that some data has been

lost during transmission (by comparing the number of events saved by the DACQ and that saved by the Harvester), the Offline Collector provides an option to re-assemble array events directly from data saved by the DACQ (see Fig. 2). The Offline Collector runs during daytime and thus adds no additional burden to runtime operations. As for the Quicklook, the conventional single-telescope analysis algorithm has been implemented and tested with data from the Whipple 10 m telescope. The implementation of a stereoscopic analysis algorithm is under way. The Harvester and the Quicklook run on a single dual-processor Xeon (2.8 GHz) server and make use of its quad-processing power via hyperthreading. The Offline Collector runs on a separate data-archiving computer.

#### 4. Standard Analysis

The standard analysis package will consist of a suite of tools (which we refer to as *htools*). Each tool is designed to perform a specific task, with HDF5 files as input and output. A set of tools may be combined in a script to accomplish more complicated tasks in a single step or to allow batch-mode processing. Standard Analysis is intended for observers to obtain more reliable results (than those from the Quicklook) from an observation, shortly after it is completed.

The modular design of *htools* makes it easy to extend the package as needed. Another important advantage is that multiple algorithms can now be accommodated within one common framework. For instance, *hclean*, for image cleaning, can take an argument to indicate whether picture/boundary cleaning or wavelet cleaning is to be applied. Similarly, different sets of optimized cut parameters can be loaded into *hcuts*, by specifying respective parameter files on the command line, and be applied to data taken in different years, with different camera configurations, or under whatever conditions which may have changed.

Besides the tools for common data analysis tasks, the package will also include general utilities for I/O, plotting, fitting, and so on. Therefore, *htools* provides a general framework, within which more sophisticated analysis algorithms may be implemented. The design of *htools* is modeled after the highly successful *f-tools* for X-ray astronomy. As the field of VHE astrophysics matures, we believe that a common data format and software platform is desirable, based on lessons learned from the early days of X-ray astronomy.

We gratefully acknowledge support from the Department of Energy, the National Science Foundation, and Enterprise-Ireland.

1. Catanese, M., & Weekes, T.C. 1999, *PASP*, 111 1193
2. Ong, R. 2003, Proc. “The Universe Viewed in Gamma-Rays”, Ed. Enomoto (Universal Academy Press: Tokyo), 25–28 Sept. 2003, Kashiwa, Japan
3. Weekes, T., et al. 2002, *APh*, 17, 221